

**HARDWARE-ASSISTED METHOD FOR THREAD SCHEDULING
USING DATA CACHE LOCALITY**

ABSTRACT OF THE DISCLOSURE

There is provided a method for scheduling threads in a multi-processor computer system having an operating system and at least one cache. In a first data structure thread ids are stored for at least some of the threads associated with a context switch performed by the operating system. Each of the thread ids uniquely identifies one of the threads. In a second data structure a plurality of entries are stored for a plurality of groups of contiguous cache lines. Each of the plurality of entries is arranged such that a thread id in the first data structure is capable of being associated with at least one of the contiguous cache lines in at least one of the plurality of groups of contiguous cache lines, the thread identified by the thread id having accessed the at least one of the contiguous cache lines in the at least one of the plurality of groups of contiguous cache lines. Patterns are mined for in the plurality of entries in the second data structure to locate multiples of a same thread id that repeat with respect to at least two of the plurality of groups of contiguous cache lines. The threads identified by the located multiples of the same thread id are mapped to at least one native thread.

The threads identified by the located multiples of the same thread may include m threads and the at least one native thread may include n threads, with m and n being integers, and m being greater than n. The threads identified by the
5 located multiples of the same thread id and any other threads identified by any other thread ids associated with the at least two of the plurality of groups of contiguous cache lines may be scheduled on the same processing unit.